arXiv:1701.08142v3 [stat.ME] 3 Mar 2017

Modelling Preference Data with the Wallenius Distribution

Clara Grazian[1] Fabrizio Leisen[2] Brunero Liseo[3]

[1] University of Oxford, U.K. [2] University of Kent, U.K.

[3] Sapienza Università di Roma, Italy

Abstract

The Wallenius distribution is a generalisation of the Hypergeometric distribution where weights are assigned to balls of different colours. This naturally defines a model for ranking categories which can be used for classification purposes. Since, in general, the resulting likelihood is not analytically available, we adopt an approximate Bayesian computational (ABC) approach for estimating the importance of the categories. We illustrate the performance of the estimation procedure on simulated datasets. Finally, we use the new model for analysing two datasets about movies ratings and Italian academic statisticians' journal preferences. The latter is a novel dataset collected by the authors.

Keywords: Approximate Bayesian Computation, Biased Urn, Movies ratings, Scientific Journals Preferences.

1 Introduction and motivations

Human beings naturally tend to rank objects in everyday life such as food, shops, singers and football teams, according to their preferences. More generally, to rank a set of objects means to arrange them in order with respect to some characteristic. Ranked data are often employed in contexts where objective and precise measurements can be impossible or unreliable and the observer collects ordinal information about preferences, judgements, relative or absolute ranking among competitors, called items. Modern technologies, such as the web, have made available a huge amount of ranked data, which can provide information about social and psychological behaviour, marketing strategies and political preferences. The codification of this information has been of interest to the statisticians since the beginning of the 20th century. The *Thurstone model* (TM) assumes that each item i is associated with a score W_i on which the comparative judgement is based; examples of unidimensional scores are the unrecorded finishing times of players in a race or any possible preference/attitude measure towards items. Item i is preferred to item j if W_i is greater than W_j, see Thurstone (1927). From the modelling point of view, this corresponds to assigning a probability $p_{ij} = \Pr(W_i > W_j)$. The *Bradley-Terry* model (BT) is a particular case

of the TM model with $p_{ij} = p_i(p_i + p_j)^{-1}$ where $p_i, p_j \geq 0$ are the item parameters reflecting the rate of each item, see Bradley and Terry (1952). Paired comparison models are always applicable to rankings after converting the latter in a suitable set of pairwise preferences. Conversely, paired comparisons of K items do not necessarily correspond to a ranking, due to the potential presence of circularites. A popular extension of the BT model is the *Plackett-Luce model* (PL). Given a set of L items and a vector of probabilities $(p_1, \ldots, p_L)$, such that $\sum_{i=1}^{L} p_i = 1$, the PL model assigns a probability distribution on all the set of possible rankings of these objects which is a function of the $(p_1, \ldots, p_K)$, see Plackett (1975) and Luce (1959). TM, BT and PL are not the only proposals in the field, and modelling ranking is an active area of research, see Marden (1995) and Alvo and Yu (2014). In this paper, we propose a new perspective on this literature by considering that rankings can be further classified into categories of different importance. Our approach makes use of an extension of the hypergeometric distribution, namely the Wallenius distribution (Wallenius, 1963).

The Wallenius distribution arises quite naturally in situations where sampling is performed without replacement and units in the population have different probabilities to be drawn. To be more specific, consider a urn with balls of c different colours: for $i = 1, \ldots, c$ there are m_i balls of colour i. In addition, colour i has a priority $\omega_i > 0$ which specifies its relative importance with respect to the other colours. A sample of n balls - with $n < \sum_{i=1}^{c} m_i$ - is drawn sequentially without replacement. The Wallenius distribution describes the probability distribution for all possible strings of balls of length n drawn from this urn. This experimental situation arises in very different contexts. For example, in auditing problems, transactions are examined by randomly selecting a single euro (or pound, or dollar) among the total amount, so larger transactions are more likely to be drawn and checked.

The Wallenius distribution was introduced by Wallenius (1963) and it is also known as the noncentral hypergeometric distribution; this alternative name is justified by the fact that, when all the priorities ω_i's are equal, one gets back to the classical hypergeometric distribution. However this name should be avoided because, as extensively discussed by Fog (2008a), this is also the name of another distribution, proposed by Fisher (1935). Although the Wallenius distribution is a very natural statistical model for the aforementioned situations, its popularity in applied settings has been prevented by the lack of a closed form expression of the probability mass function: see Section 2 for details.

The gist of this paper is the use of the priorities vector $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_c)$ of the Wallenius distribution as a measure of importance for different values of a categorical variable.

In particular, we analyse two datasets, where we aim at ranking the categories rather than the items. The first dataset considers data downloaded from the MovieLens website, which consists of 105,339 ratings across 10,329 movies performed by 668 users. In this framework, it is of interest to classify the different genres in terms of satisfaction, in order to provide some useful feedback to users and/or providers.

The second dataset considers data we collected between October and November 2016 among Italian academic statisticians. They indicated their journal preferences

from the 2015 ISI "Statistics and Probability" list of Journals. In this context, we are interested in ranking the journal categories in order to provide a description of the research interests of the Italian Statistical community.

We adopt a Bayesian methodology which allows us to overcome the computational problems related to the lack of a closed form expression of the probability mass function of the Wallenius distribution. We propose a novel approximate Bayesian computational approach (Marin et al., 2012), where the vector of summary statistics is represented by the relative frequencies of the different categories and the acceptance mechanism is based on the distance in variation (Bremaud, 1998)

The paper is organized as follows: in Section 2 we will present the Wallenius distribution, in Section 3 our approximated inferential approach will be described, based on an ABC algorithm. The performance of the algorithm will be tested in several examples, first in an extensive simulation study (Section 4) and then on real datasets (Section 5). A discussion concludes the paper.

2 The Wallenius Distribution

Consider a urn with N balls of c different colours. There are m_i balls of the i-th colour, so that $\sum_i^c m_i = N$. In this situation, the multivariate Hypergeometric distribution is the discrete probability distribution which describes the sampling without replacement of n balls. In this framework, the probability of drawing a ball of a certain colour is proportional to the number of balls of the same colour. It is possible to generalise the experiment with a biased sampling of balls. For instance, each colour may have a different priority or importance, say $\omega_i > 0$, $i = 1, \ldots, c$. Suppose we have drawn n balls without replacement from the urn and let $\boldsymbol{X}_n = (X_{1n}, X_{2n}, \ldots, X_{cn})$ denote the frequencies of balls of different colours in the sample. Let Z_n be the colour of the ball drawn at time n. In this setting, the probability that the next ball is of colour i could also depend on its priority and is defined as

$$P\left(Z_{n+1} = i | \boldsymbol{X}_n\right) = \frac{\left(m_i - X_{in}\right)\omega_i}{\sum_{j=1}^{c}\left(m_j - X_{jn}\right)\omega_j}. \tag{1}$$

Wallenius (1963) provided the above expression and the probability mass function of $\boldsymbol{X}_n$ for the case $c = 2$. Chesson (1976) derived the following general expression. For a given integer n, and parameters $\boldsymbol{m} = (m_1, \ldots, m_c)$ and $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_c)$, the probability of observing a vector of colour frequencies $\boldsymbol{x} = (x_1, \ldots, x_c)$ is

$$P(\boldsymbol{x}; n, \boldsymbol{m}, \boldsymbol{\omega}) = \prod_{j=1}^{c}\binom{m_j}{x_j}\int_0^1 \prod_{j=1}^{c}\left(1 - t^{\omega_j/d}\right)^{x_j} dt, \tag{2}$$

where $\sum_{i=1}^{c} x_i = n$ and $d = \sum_{j=1}^{c}\omega_j(m_j - x_j)$. When $\omega_i = \omega$, for every $i = 1, \ldots, c$, the Wallenius distribution reduces to the multivariate Hypergeometric distribution. This can be easily shown by considering, without loss of generality, $\omega = 1$ and $c = 2$.

In this particular case, the probability mass function simplifies to

$$P(x;n,m) = \binom{m}{x}\binom{N-m}{n-x}\int_0^1 \left(1-t^{1/d}\right)^n.$$

The change of variable $z = t^{1/d}$ leads to

$$\begin{aligned} P(x;n,m) &= \binom{m}{x}\binom{N-m}{n-x} d\int_0^1 (1-z)^n z^{d-1} dz \\ &= \binom{m}{x}\binom{N-m}{n-x}\frac{\Gamma(d+1)\Gamma(n+1)}{\Gamma(n+d+1)}. \end{aligned}$$

By observing that $d = N - n$, the probability mass function reduces to

$$P(x;n,m) = \binom{m}{x}\binom{N-m}{n-x}\Big/\binom{N}{n},$$

which is the probability mass function of the Hypergeometric distribution when two colours are considered.

The Wallenius distribution has been underemployed in the statistical literature mainly because the integral appearing in (2) cannot be solved in a closed form and numerical approximations are necessary. Fog (2008a) has made a substantial contributions in this direction, providing approximations based either on asymptotic expansions or numerical integration. To our knowledge, the Wallenius distribution has only been used in a limited number of applications, mainly devoted to auditing problems (Gillett, 2000), ecology (Manly, 1974), vaccine efficacy (Hernández-Suárez and Castillo-Chavez, 2000) and modeling of RNA sequences (Gao et al., 2011). In this work, we propose a new use of the Wallenius distribution to rank the categories of a discrete variable based on preferences. This is motivated by the sampling nature of the Wallenius distribution where an importance ω_j is associated with category j. The highest ω_j's represent the most popular categories. This naturally defines a new model which allows us to rank preferences. Recently, the development of social networks and the competitive pressure to provide customized services has motivated many new ranking problems involving hundreds or thousands of objects. Recommendations on products such as movies, books and songs are typical examples in which the number of objects is extraordinarily large. In recent years, many researchers in statistics and computer science have developed models to handle such big data. For instance, in Section 5 we consider the problem of ranking customer movie choices in terms of genres such as Comedy, Drama and Science Fiction. We consider data dowloaded from the MovieLens website (`www.grouplens.org`) which consists of 105,339 online ratings of 10,329 movies by 668 raters on a scale of 1-5. We rank the categories by estimating the priority parameters of the Wallenius distribution by using an approximate Bayesian approach. In particular, the next section introduces an algorithm which allows us to deal with the integral in equation (2).

3 Bayesian Inference for the Wallenius distribution

Let $\boldsymbol{x}_h = (x_{h1}, \ldots, x_{hc})$ be a draw of n_h balls from the Wallenius urn described in equation (2), where $h = 1, \ldots, k$ and $\sum_{j=1}^{c} x_{hj} = n_h$. In this paper we adopt a Bayesian approach where the vector of parameters $\boldsymbol{\omega}$ is unknown. For a given prior distribution $\pi(\boldsymbol{\omega})$ the posterior distribution is

$$\pi(\boldsymbol{\omega}|\boldsymbol{x}_1, \ldots, \boldsymbol{x}_k) \propto \pi(\boldsymbol{\omega}) \prod_{h=1}^{k} \left[\int_0^1 \prod_{j=1}^{c} \left(1 - t_h^{\omega_j/d_h}\right)^{x_{hj}} dt_h \right], \tag{3}$$

with $d_h = \sum_{j=1}^{c} \omega_j (m_j - x_{hj})$. The above posterior distribution depends on k different integrals which cannot be reduced to a closed form. This makes the implementation of standard Markov Chain Monte Carlo (MCMC) methods for estimating $\boldsymbol{\omega}$ difficult. Indeed, most MCMC methods rely on the direct evalutation of the unnormalized posterior distribution (3). Although there are many available routines, in different software packages, to evaluate univariate integrals, we noticed that they lack accuracy especially for large values of n and $\boldsymbol{m}$. We believe that this problem has had a strong negative impact on the popularization of the Wallenius distribution despite a need for interpretable models in the applied setting. For instance, the Wallenius distribution arises naturally in genetics as an alternative to the Fisher exact test, see Gao et al. (2011) and the references therein.

In this section, we propose an algorithm which allows to sample from the posterior distribution introduced in (3). The algorithm belongs to the class of approximate Bayesian computational (ABC) methods. This aapproach is philosophically different from the standard MCMC methods since their implementation only require to draw samples from the generating model for a given parameter value. In the case of the Wallenius distribution, the task of generating draws is not hard, making the use of ABC particularly straightforward. Fog (2008b) provided methods and algorithms to sample from the Wallenius distribution. He also made available a reliable R package, called `BiasedUrn`, which has been used extensively in this work.

The ABC methodology can be considered as a (class of) popular algorithms that achieves posterior simulation by avoiding the computation of the likelihood function: see Beaumont (2010) or Marin et al. (2012) for recent surveys. As remarked by Marin et al. (2012), the first genuine ABC algorithm was introduced by Pritchard et al. (1999) in a population genetics setting. Explicitly, we consider a parametric model $\{f(\cdot \mid \theta), \theta \in \Theta\}$ and suppose that a dataset $\boldsymbol{y} \in \mathcal{D} \subset \mathbb{R}^n$ is observed. Let $\varepsilon > 0$ be a tolerance level, η a summary statistic (which often is not sufficient) defined on $\mathcal{D}$ and ρ a distance or metric on $\eta(\mathcal{D})$. Let π be a prior distribution for θ, the ABC algorithm is described in Algorithm 1. The basic idea behind the ABC is that, for a small (enough) ε and a representative summary statistic, we can obtain a reasonable approximation of the posterior distribution. The practical implementation of an ABC algorithm requires the selection of a suitable summary statistic, a distance and a tolerance level. In our

Algorithm 1 ABC Rejection algorithm

1: **for** $l = 1, \cdots, T$ **do**
2: **repeat**
3: Generate θ' from the prior distribution $\pi(\cdot)$
4: Generate z from the likelihood $f(\cdot \mid \theta')$
5: **until** $\rho(\eta(\mathbf{z}), \eta(\mathbf{y})) < \varepsilon$
6: Set $\theta_l = \theta'$
7: **end for**

specific case we summarized the data by using the arithmetic mean of the observed and simulated frequency vectors, i.e., at the ℓ-th iteration of pseudo data generation, we have

$$\eta(\boldsymbol{x}^{(\ell)}) = \widehat{\boldsymbol{p}}^{(\ell)} = \frac{1}{k}\sum_{h=1}^{k} \boldsymbol{p}_h^{(\ell)}, \tag{4}$$

with

$$\boldsymbol{p}_h^{(\ell)} = \left(\frac{x_{h1}^{(\ell)}}{n_h}, \ldots, \frac{x_{hc}^{(\ell)}}{n_h}\right)$$

to be compared with the relative frequencies observed in the sample

$$\eta(\boldsymbol{x}^{(t)}) = \widehat{\boldsymbol{p}}^{(t)} = \frac{1}{k}\sum_{h=1}^{k} \boldsymbol{p}_h^{(t)}.$$

with

$$\boldsymbol{p}_h^{(t)} = \left(\frac{x_{h1}}{n_h}, \ldots, \frac{x_{hc}}{n_h}\right).$$

Since the frequencies $\widehat{\boldsymbol{p}}^{(\ell)} = \left(\widehat{p}_1^{(\ell)}, \ldots, \widehat{p}_c^{(\ell)}\right)$ and $\widehat{\boldsymbol{p}}^{(t)} = (\widehat{p}_1, \ldots, \widehat{p}_c)$ can be interpreted as discrete probability distributions, it is natural to use the distance in variation (Bremaud, 1998),

$$\rho(\widehat{\boldsymbol{p}}^{(\ell)}, \widehat{\boldsymbol{p}}^{(t)}) = \frac{1}{2}\sum_{j=1}^{c} \left|\widehat{p}_j^{(\ell)} - \widehat{p}_j\right| \tag{5}$$

For the setting of the tolerance level we refer to the Section 4 where the algorithm will be tested on simulated data.

The prior distribution

The vector of parameters $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_c)$ assumes values in $\mathbb{R}_+^c$ and different priors can be considered. However, one must take into account that the priority parameters ω_j must be interpreted in a relative way. Therefore, a re-parametrization of them in terms of a normalization factor seems more appropriate. For instance we can normalize

the priority parameters in terms of their sum and this provides an interpretation in terms of weights of importance. In this perspective, a reasonable prior for the weights is the product of independent and identical Gamma distributions $\mathcal{G}a(\psi\xi, \xi)$, which is the choice we will use throughout the paper, although a different elicitation of the prior weights can be easily accomodated.

Alternative computational approaches

The R package `BiasedUrn` allows the approximate numerical evaluation of the probability mass function of the Wallenius distribution. In a classical setting this makes feasible the computation of the MLE. In a Bayesian setting this enables the implementation of standard MCMC algorithms, such as the Metropolis-Hastings sampler. Nonetheless, we deem more appropriate to use the ABC approach illustrated in this section. First, the output of the Bayesian approach is far richer than the one available in a classical setting. For instance, in Section 5.2 we are able to easily compute the probability $p_{ij} = \Pr(\omega_i > \omega_j)$. Second, standard MCMC methods require repeated evaluations of the likelihood function. This could lead to an unsustainable computational burden compared to ABC.

Identifiability

We conclude this section with a cautionary remark abouth parameter identifiability. Consider the p.m.f. of the Wallenius random vector displayed in (2). Given $\mathbf{x}$ and $\mathbf{m}$, the p.m.f. is exactly equal when we consider two different vectors $\boldsymbol{\omega\prime}$ and $\boldsymbol{\omega}$ such that $\boldsymbol{\omega\prime} = \kappa\boldsymbol{\omega}$ for $\kappa > 0$. This implies an identifiability issue which can be tackled by normalizing the vector $\boldsymbol{\omega}$ with respect to one component, constraining the estimation to the remaining $(c-1)$ components of $\boldsymbol{\omega}$. In the examples in the following sections, we normalized, in general, with respect to the last value ω_c, the only exception being the third simulation study.

4 Simulation Study

In order to test Algorithm 1 with the summary statistics shown in Section 3, we have conducted an extensive simulation study, with different scenarios. We performed $1,000$ repeated simulations of k draws from the Wallenius distribution where each draw consists of a number n_h $(h = 1, \cdots, k)$ of balls. We use the prior distribution defined in Section 3, i.e. independent gamma priors $\mathcal{G}a(\psi\xi, \xi)$, such that the expected value of each ω_i is a fixed value ψ and the variance is ψ/ξ, where ξ is a hyperparameter with distribution $\mathcal{E}xp(1)$. As already stated in Section 3, we use the summary statistics and the distance in variation defined in equations (4) and (5). The tolerance level ε has been chosen with a pilot simulation in order to produce a sample of size 10^4. We fixed the number of colours equal to $c = 3$ and considered three different configurations of the other parameters of the Wallenius distribution:

- for the number of balls for each colour, we use $\mathbf{m} = (15, 8, 7)$, $\mathbf{m} = (10, 10, 10)$ and $\mathbf{m} = (20, 5, 5)$;
- for the importance weight for each colour, we use $\boldsymbol{\omega} = (5, 2, 1)$, $\boldsymbol{\omega} = (1, 1, 1)$ and $\boldsymbol{\omega} = (1, 2, 5)$;
- for the number of observations in the sample, we consider $k = 5$, $k = 50$ and $k = 1000$.

In the first two simulation experiments, we have normalised the importance weights with respect to the last value ω_3, which is in both cases set equal to 1; on the other hand, we normalised with respect to the first value ω_1 in the third experiment. The value of n_h has been simulated from a discrete uniform distribution which takes values in $[5, 30]$. The results are available in Tables 1, 2 and 3. As expected, for all the configurations considered, as the sample size k increases, the results improve in terms of accuracy. On the other hand, the results are less accurate for configurations where $\boldsymbol{\omega}$ and $\mathbf{m}$ follow an opposite ordering. For instance, when $\omega_1 < \omega_2 < \omega_3$ and $m_1 > m_2 > m_3$ or $\omega_1 > \omega_2 > \omega_3$ and $m_1 < m_2 < m_3$. The case of equal importance weights, which corresponds to the classical multivariate Hypergeometric distribution, is well estimated with our approach. We have performed simulations with a higher number of balls and/or a higher number of colours (not reported here), with similar results.

5 Real Data Applications

We now apply the proposed approach to two real datasets, in order to assess the applicability and the performance of the algorithm. In both cases, we obtain the ratings of a group of individuals about specific elements from a list. Each individual may choose the number of elements to rate. The elements are then grouped in categories and the goal is to provide a ranking of the categories. By using the urn terminology of Section 2, the categories are the colours and each element from the list is a ball; the aim of the analysis is to perform inference on the importance weights of each colour.

5.1 Movies dataset

This dataset describes 5-star (with half-star increments) rating from MovieLens, a movie recommendation service (http://grouplens.org/datasets/movielens/). The dataset may change over time. We consider the dataset which contains 105,339 ratings across 10,329 movies. These data were created by 668 users between April 03, 1996 and January 09, 2016. This dataset was generated on January 11, 2016. Users were randomly selected by MovieLens, with no demographic information, and each of them has rated at least 20 movies. The movies in the dataset were described by genre, following the *IMDb* information (https://www.themoviedb.org/); eighteen genres were considered, plus an empty category. Each film may be described by more than one genre. In this case we have proceeded as follows: we have ordered the genres in terms

Table 1: Simulation study with true $\boldsymbol{\omega} = (5, 2, 1)$ ($k = 5$, $k = 50$, $k = 1000$). Each row represents the average of the posterior medians computed over $1,000$ replications of the experiment (the standard deviations are in brackets). We have renormalized the simulated values by fixing the third element equal to one.

	$\mathbf{m} = (\mathbf{15}, \mathbf{8}, \mathbf{7})$	$\mathbf{m} = (\mathbf{10}, \mathbf{10}, \mathbf{10})$	$\mathbf{m} = (\mathbf{7}, \mathbf{8}, \mathbf{15})$
k=5	5.60, 3.96, 1.00	5.13, 3.25, 1.00	5.85, 4.42, 1.00
	(2.51, 2.63, -)	*(2.19, 1.88, -)*	*(2.52, 2.96, -)*
k=50	5.10, 3.11, 1.00	4.95, 2.65, 1.00	5.51, 3.99, 1.00
	(0.77, 0.63, -)	*(0.73, 0.40, -)*	*(0.89, 0.98, -)*
k=1000	5.01, 2.01, 1.00	4.91, 2.61, 1.00	5.38, 3.97, 1.00
	(0.20, 0.09, -)	*(0.18, 0.09, -)*	*(0.23, 0.31, -)*

Table 2: Simulation study with true $\boldsymbol{\omega} = (1, 1, 1)$ ($k = 5$, $k = 50$, $k = 1000$). Each row represents the average of the posterior medians computed over $1,000$ replications of the experiment (the standard deviations are in brackets). We have renormalized the simulated values by fixing the third element equal to one.

	$\mathbf{m} = (\mathbf{15}, \mathbf{8}, \mathbf{7})$	$\mathbf{m} = (\mathbf{10}, \mathbf{10}, \mathbf{10})$	$\mathbf{m} = (\mathbf{7}, \mathbf{8}, \mathbf{15})$
k=5	1.01, 1.37, 1.00	0.99, 1.30, 1.00	1.04, 1.49, 1.00
	(0.45, 0.81, -)	*(0.42, 0.66, -)*	*(0.58, 1.27, -)*
k=50	1.03, 1.23, 1.00	1.00, 1.19, 1.00	1.10, 1.36, 1.00
	(0.12, 0.17, -)	*(0.12, 0.15, -)*	*(0.14, 0.22, -)*
k=1000	1.05, 1.01, 1.00	1.00 1.17, 1.00	1.22, 1.34, 1.00
	(0.03, 0.03, -)	*(0.03, 0.03, -)*	*(0.03, 0.06, -)*

Table 3: Simulation study with true $\boldsymbol{\omega} = (1, 2, 5)$ ($k = 5$, $k = 50$, $k = 1000$). Each row represents the average of the posterior medians computed over $1,000$ replications of the experiment (the standard deviations are in brackets). We have renormalized the simulated values by fixing the first element equal to one.

	$\mathbf{m} = (\mathbf{15}, \mathbf{8}, \mathbf{7})$	$\mathbf{m} = (\mathbf{10}, \mathbf{10}, \mathbf{10})$	$\mathbf{m} = (\mathbf{7}, \mathbf{8}, \mathbf{15})$
k=5	1.00, 3.39, 5. 03	1.00, 3.38, 5.31	1.00, 4.12, 4.88
	(-, 2.15, 2.18)	*(-, 2.08, 2.40)*	*(-, 2.91, 2.14)*
k=50	1.00, 2.58, 4.88	1.00, 2.67, 4.95	1.00, 2.79, 4.76
	(-, 0.36, 0.72)	*(-, 0.41, 0.75)*	*(-, 0.45, 0.72)*
k=1000	1.00, 2.53, 4.84	1.00, 2.61, 4.91	1.00, 2.70, 4.70
	(-, 0.08, 0.17)	*(-, 0.09, 0.17)*	*(-, 0.10, 0.18)*

of their generality and then assigned to the movie the least general genre with which it was described. We have decided the following order (from the less general to the most general): Animation → Children → Musical → Documentary → Horror → Sci-Fi → Film Noir → Crime → Fantasy → War → Western → Mistery → Action → Thriller → Adventure → Romance → Comedy → Drama. Of course, this is an experimental choice, which may affect the results. Since the movies can be cross-classified, an interesting (and more realistic) development would be considering a model which can take into account this feature; this is left for further research. We have then replicated the same prior choice and the same choices of distance and vector of summary statistics described in Section 4. The tolerance level ε has been chosen with a pilot simulation in order to produce a sample of size 10^5. In this particular case, we have used $\varepsilon = 0.25$. Table 5.1 displays the posterior mean estimates of the vector of importance weights $\boldsymbol{\omega}$. The importance weights seem to be very close, with small differences among them. This suggests that there is not a category which is particularly popular. Nonetheless, we can observe a slightly preference for the *War* and *Crime* genres and less interest in the *Science fiction* and *Action* genres. We believe that this similarity in the importance weights is due to an excessive number of categories in the movies dataset. In this setting the graphical comparison of the marginal posterior distributions can provide a better insight on the customer preferences. Figure 1 shows that there is more variability in the users preferences to choose a particular movie genre, such as *War* or *Mistery*.

Category	$\boldsymbol{\omega}$	**Category**	$\boldsymbol{\omega}$	**Category**	$\boldsymbol{\omega}$
War	1.969	*Western*	1.616	*Adventure*	1.566
	(1.620)		(1.862)		(1.467)
Crime	1.762	*Animation*	1.601	*Romance*	1.496
	(1.078)		(1.150)		(0.917)
Children	1.667	*Drama*	1.596	*Horror*	1.420
	(2.021)		(1.082)		(1.478)
Film-Noir	1.661	*Documentary*	1.594	*Comedy*	1.277
	(1.803)		(1.759)		(0.926)
Thriller	1.652	*Fantasy*	1.591	*Sci-Fi*	1.264
	(1.556)		(1.342)		(0.731)
Musical	1.632	*Mystery*	1.576	*Action*	1.000
	(1.591)		(1.158)		-

Table 4: Posterior mean estimates and standard deviations (in brackets) of the vector of importance weights $\boldsymbol{\omega}$ for each genre with tolerance level $\varepsilon = 0.25$. We have fixed the weight for the *Action* genre equal to 1.

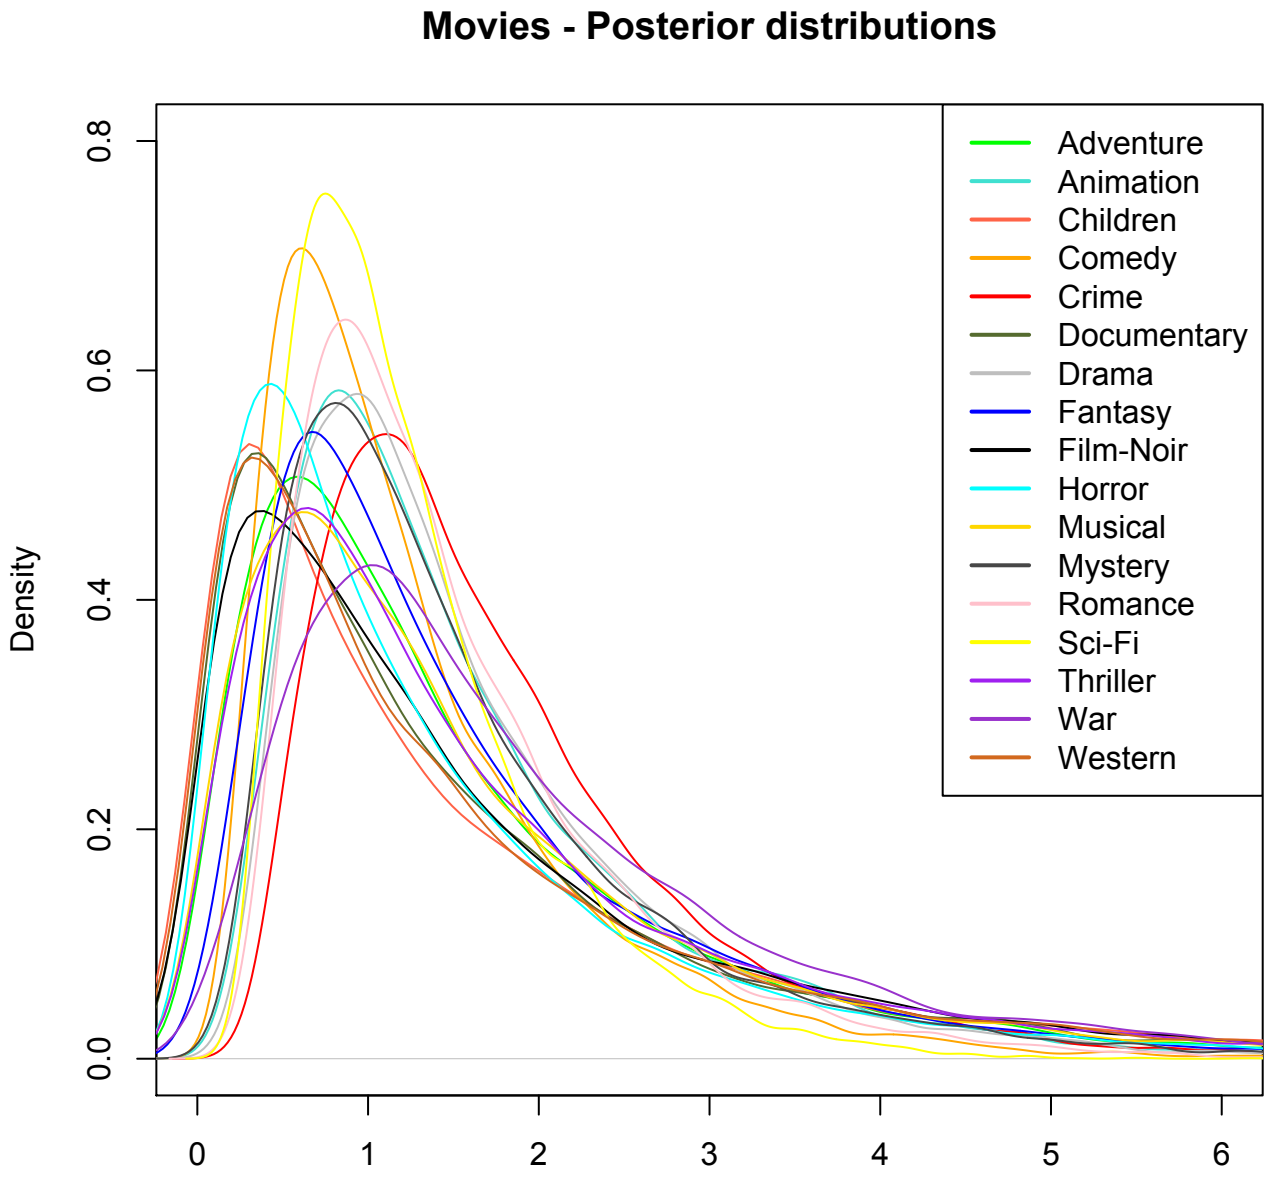

Figure 1: Approximations of the posterior distributions of the weights $\boldsymbol{\omega}$ for each category included in the Movies dataset. We have fixed the weight for the *Action* genre equal to 1.

5.2 Statistical Journals dataset

The scientific areas (or "settori scientifici disciplinari", S.S.D.) are a characterization used in the academic Italian system to classify knowledge in higher education. The sectors are determined by the Italian Ministry of Education. In particular, there are 367 S.S.D., divided into 14 macro-areas and each member of the academic staff pertains to a single sector. We have performed a survey on the preferences of the researchers in Statistics (Sector SECS-S/01) of Italian universities about the available scientific journals. It should be noted that researchers in Probability and Mathematical Statistics, Medical, Economic and Social Statistics are not included in this survey, because they pertain to different sectors. We have considered only staff with both teaching and research contracts. Postdoctoral fellows and PhD students have been excluded. In this survey we have used the 2015 "Statistics and Probability" list of journals of the Institute for Scientific Information (ISI). We have asked to SESC-S/01 researchers to indicate their preferences in this list, between a minimum of ten and a maximum of twenty. One difference from the Movies example of Section 5.1 is that the participants do not have to indicate the level of their preference, only a list of journals which each of the participants considers either

- prestigious and/or
- likely for a potential submission and/or
- professionally significant (in terms of frequency of readings).

The survey was conducted between 25th October 2016 and 4th November 2016. We have collected 174 responses, distributed, in terms of role, as follows: 49 Full professors (Professori Ordinari), 72 Associate Professors (Professori Associati) and 53 Assistant Professors, both fixed-term and tenure-track (Ricercatori a tempo indeterminato e a tempo determinato). We have then grouped the journals by category, considering five main classes of interest: *Methodology, Probability, Applied Statistics, Computational Statistics* and *Econometrics and Finance*. The list of journals and relative category is available in the Appendix. Among the 124 journals available in the "Statistics and Probability" ISI list, we have classified 23 journals in *Probability*, 45 in *Methodology*, 34 in *Applied Statistics*, 9 in *Computational Statistics* and 13 in *Econometrics and Finance*. We assume the Wallenius distribution for modelling the dataset, where c represents the number of the categories. The preferences of each respondent are summarized in a vector where the position of each entry represents the number of journals falling in the corresponding category. We consider that this vector is a realization of the Wallenius distribution.

The results are available in Figure 2, Figure 3 and Table 6, which show that there seems to be a preference for the research in Methodological and Computational Statistics among the researchers in Statistics and less interest in journals of Probability. As already stated, this should highlight the fact that researchers in Mathematical Statistics and Probability do not pertain to the investigated sector. The results also show that the effect of a decrease of the tolerance level seems to be a concentration of the posterior distributions of the importance weights $\boldsymbol{\omega}$, except for the weight relative

Table 5: Each entry p_{ij} of the matrix represents the ABC approximation of $\Pr(\omega_i > \omega_j)$.

	ω_1	ω_2	ω_3	ω_4	ω_5
ω_1	-	1.000	0.999	0.394	1.000
ω_2		-	0.000	0.000	0.226
ω_3			-	0.104	0.951
ω_4				-	0.992
ω_5					-

Table 6: Posterior mean estimates and standard deviations (in brackets) of the vector of importance weights $\boldsymbol{\omega}$ for each category of journals and for different tolerance levels. We have fixed the weight for the *Probability* category equal to 1.

	Methodology	**Applied**	**Computational**	**Econometrics**
$\boldsymbol{\omega}$	4.821	3.251	3.615	1.751
$\varepsilon = 0.130$	*(2.429)*	*(1.768)*	*(2.664)*	*(1.421)*
$\boldsymbol{\omega}$	6.183	4.165	6.245	1.994
$\varepsilon = 0.085$	*(2.916)*	*(2.505)*	*(3.383)*	*(1.545)*
$\boldsymbol{\omega}$	6.758	4.558	7.326	2.095
$\varepsilon = 0.070$	*(3.173)*	*(2.192)*	*(3.723)*	*(1.512)*

to the Computational journals, for which there is a shift. As an explanation of this, one should consider that this category is under-represented in the list (at least, with our classification) with respect to the others. Table 5 shows the estimated pair comparison probabilities for the journal categories.

6 Conclusions

In this paper we considered the problem of ranking categories. We proposed a novel model based on the Wallenius distribution. In terms of an urn scheme, it represents a generalization of the hypergeometric distribution with an extra-vector of parameters $\boldsymbol{\omega}$, which represents the importance of the different types of balls in the urn. So far the Wallenius model has been definitely under-employed, due to the analytical intractability of the probability mass function. In this work we proposed an approximate Bayesian computational algorithm which provides a fast and reliable approach to the estimation of the vector of priorities $\boldsymbol{\omega}$. Our method is easy to implement and it might be very useful in several statistical applications where balls are drawn from the urn in a biased fashion. Paradigmatic examples of the importance of the Wallenius model especially appear in auditing where transactions are randomly checked with probability proportional to their monetary value. We analysed two datasets concerning movies ratings and Italian academic statisticians' journal preferences. The ABC algorithm allows us to estimate the importance of movies

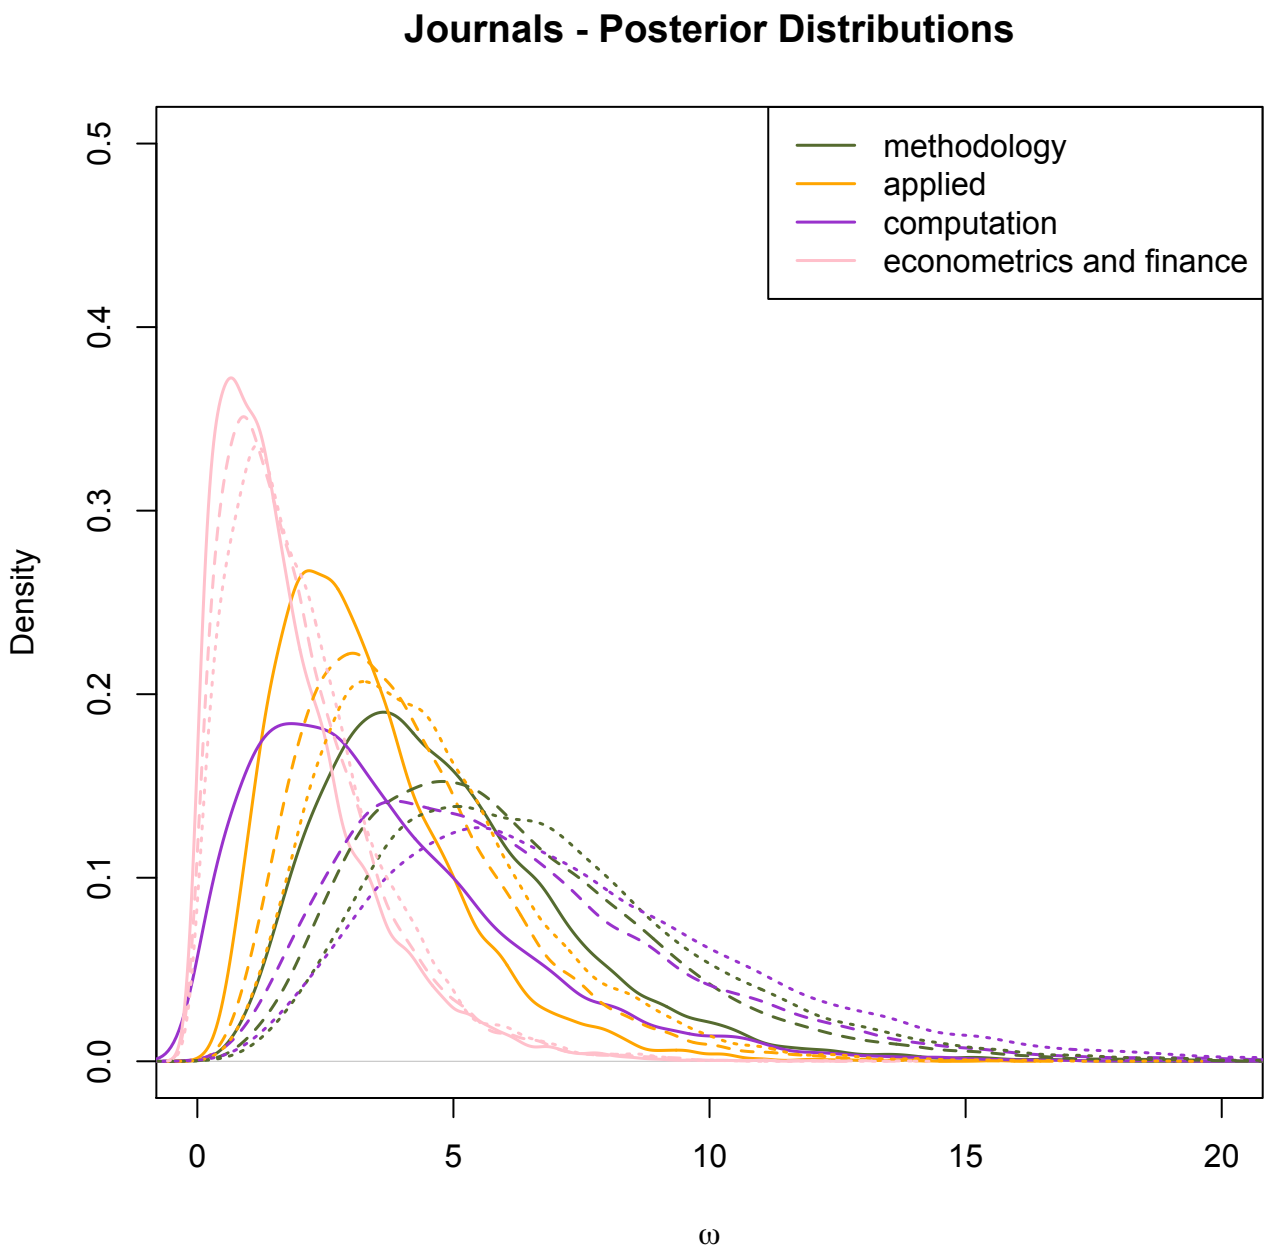

Figure 2: Approximations of the posterior distributions of the weigths $\boldsymbol{\omega}$ for each category included in the Journals dataset. We have fixed the weight for the *Probability* category equal to 1. Solid lines represent the approximations for $\varepsilon = 0.130$, dashed lines for $\varepsilon = 0.085$ and dotted lines for $\varepsilon = 0.070$.

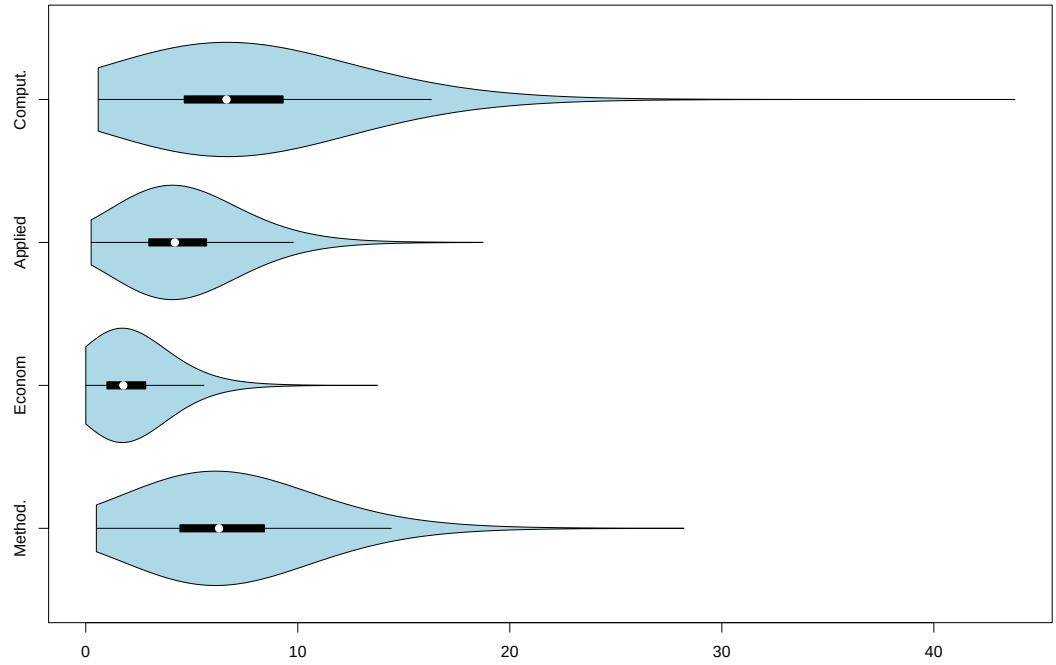

Figure 3: Violin plots of the posterior distributions of the weigths $\boldsymbol{\omega}$ for each category included in the Journals dataset with $\varepsilon = 0.070$. We have fixed the weight for the *Probability* category equal to 1.

categories or journal preferences under the assumption of a Wallenius generating model. Future work will focus on the use of the Wallenius distribution to other areas of application and on the estimation of the category multiplicities $\mathbf{m}$ given the knowledge of the importance weights $\boldsymbol{\omega}$.

Acknowledgements

The authors are very grateful to Martin Ridout for his valuable comments on a first draft of the paper. This project has been funded by the Royal Society International Exchanges Grant "Empirical and Bootstrap Likelihood Procedures for Approximate Bayesian Inference". Fabrizio Leisen was supported by the European Community's Seventh Framework Programme [FP7/2007-2013] under grant agreement no: 630677.

References

Alvo, M. and Yu, P. L. (2014). *Statistical Methods for Ranking Data*. Springer, New York.

Beaumont, M. (2010). Approximate Bayesian computation in evolution and ecology. *Annual Review of Ecology, Evolution, and Systematics* **41**, 379–406.

Berger, J. O., Bernardo, J. M. and Sun, D. (2015). Overall objective priors. *Bayesian Analysis* **10**, 189–221.

Bradley, R. A. and Terry, M. E. (1952). Rank analysis of incomplete block designs. I: The method of paired comparisons. *Biometrika* **39**, 324–345.

Bremaud, P. (1998). *Markov chains: Gibbs fields, Monte Carlo simulation and queues*. Springer-Verlag: New York.

Chesson, J. (1976). A non-central multivariate hypergeometric distribution arising from biased sampling with application to selective predation. *J. Appl. Probab.* **13**, 795–797.

Fisher, R. (1935). The logic of inductive inference. *J. Roy. Statist. Soc.* **98**, 39–82.

Fog, A. (2008a). Calculation Methods for Wallenius' Noncentral Hypergeometric Distribution. *Communications in Statistics - Simulation and Computation* **37**, 258–273.

Fog, A. (2008b). Sampling Methods for Wallenius' and Fisher's Noncentral Hypergeometric Distributions. *Communications in Statistics - Simulation and Computation* **37**, 241–257.

Gao, L., Fang, Z., Zhang, K., Zhi, D. and Cui, X. (2011). Length bias correction for RNA-seq data in gene set analyses. *Bioinformatics* **27**, 662–669.

Gillett, P. R. (2000). Monetary unit sampling: a belief-function implementation for audit and accounting applications. *International Journal of Approximate Reasoning* **25**, 43–70.

Hernández-Suárez, C. M. and Castillo-Chavez, C. (2000). Urn models and vaccine efficacy. *Statistics in Medicine* **19**.

Luce, R. D. (1959). *Individual choice behavior: A theoretical analysis*. John Wiley and Sons Inc., New York.

Manly, B. J. (1974). A Model for Certain Types of Selection Experiments. *Biometrics* **30(2)**, 281–294.

Marden, J. (1995). *Analyzing and modeling rank data*. Chapman and Hall, London.

Marin, J. M., Robert, C. P. and Pudlo, P. (2012). Approximate Bayesian computational methods. *Statistics and Computing* **22**, 1167–1180.

Plackett, R. L. (1975). The analysis of permutations. *J. Roy. Statist. Soc. Ser. C Appl. Statist.* **24**, 193—202.

Pritchard, J., Seielstad, M., Perez-Lezaun, A. and Feldman, M. (1999). Population growth of human Y chromosomes: a study of Y cromosom microsatellites. *Molecular Biology and Evolution* **16**, 1791–1798.

Thurstone, L. L. (1927). A Law of Comparative Judgment. *Psychological Review* **34**, 273–286.

Wallenius, K. T. (1963). *Biased Sampling: The Non-Central Hypergeometric Probability Distribution - Department of Statistics - Stanford University*. Ph.D. thesis, Department of Statistics - Stanford University.

A Appendix

Table A.1: Journals in the Probability category

Probability
ADVANCES IN APPLIED PROBABILITY
ANNALES DE L INSTITUT HENRI POINCARE - PROBABILITES ET STATISTIQUES
ANNALS OF APPLIED PROBABILITY
ANNALS OF PROBABILITY
COMBINATORICS PROBABILITY and COMPUTING
ELECTRONIC COMMUNICATIONS IN PROBABILITY
ELECTRONIC JOURNAL OF PROBABILITY
INFINITE DIMENSIONAL ANALYSIS QUANTUM PROBABILITY AND RELATED TOPICS
JOURNAL OF APPLIED PROBABILITY
JOURNAL OF THEORETICAL PROBABILITY
MARKOV PROCESSES AND RELATED FIELDS
METHODOLOGY AND COMPUTING IN APPLIED PROBABILITY
PROBABILITY AND MATHEMATICAL STATISTICS-POLAND
PROBABILITY IN THE ENGINEERING AND INFORMATIONAL SCIENCES
PROBABILITY THEORY AND RELATED FIELDS
RANDOM MATRICES-THEORY AND APPLICATIONS
STOCHASTIC ANALYSIS AND APPLICATIONS
STOCHASTIC MODELS
STOCHASTIC PROCESSES AND THEIR APPLICATIONS
STOCHASTICS AND DYNAMICS
STOCHASTICS-AN INTERNATIONAL JOURNAL OF PROBABILITY AND STOCHASTIC REPORTS
THEORY OF PROBABILITY AND ITS APPLICATIONS
UTILITAS MATHEMATICA

Table A.2: Journals in the Methodology category

Methodology
ADVANCES IN DATA ANALYSIS AND CLASSIFICATION
ALEA-LATIN AMERICAN JOURNAL OF PROBABILITY AND MATHEMATICAL STATISTICS
AMERICAN STATISTICIAN
ANNALS OF STATISTICS
ANNALS OF THE INSTITUTE OF STATISTICAL MATHEMATICS
ANNUAL REVIEW OF STATISTICS AND ITS APPLICATION
ASTA-ADVANCES IN STATISTICAL ANALYSIS
AUSTRALIAN and NEW ZEALAND JOURNAL OF STATISTICS
BAYESIAN ANALYSIS
BERNOULLI
BIOMETRIKA
BRAZILIAN JOURNAL OF PROBABILITY AND STATISTICS
CANADIAN JOURNAL OF STATISTICS-REVUE CANADIENNE DE STATISTIQUE
COMMUNICATIONS IN STATISTICS-THEORY AND METHODS
ELECTRONIC JOURNAL OF STATISTICS
ESAIM-PROBABILITY AND STATISTICS
EXTREMES
FUZZY SETS AND SYSTEMS
HACETTEPE JOURNAL OF MATHEMATICS AND STATISTICS
INTERNATIONAL JOURNAL OF GAME THEORY
INTERNATIONAL STATISTICAL REVIEW
JOURNAL OF MULTIVARIATE ANALYSIS
JOURNAL OF NONPARAMETRIC STATISTICS
JOURNAL OF STATISTICAL PLANNING AND INFERENCE
JOURNAL OF THE AMERICAN STATISTICAL ASSOCIATION
JOURNAL OF THE KOREAN STATISTICAL SOCIETY
JOURNAL OF THE ROYAL STATISTICAL SOCIETY SERIES B STATISTICAL METHODOLOGY
JOURNAL OF TIME SERIES ANALYSIS
LIFETIME DATA ANALYSIS
METRIKA
REVSTAT-STATISTICAL JOURNAL
SCANDINAVIAN JOURNAL OF STATISTICS
SEQUENTIAL ANALYSIS-DESIGN METHODS AND APPLICATIONS
SPATIAL STATISTICS
STATISTICA NEERLANDICA
STATISTICA SINICA
STATISTICAL ANALYSIS AND DATA MINING
STATISTICAL METHODOLOGY
STATISTICAL METHODS AND APPLICATIONS
STATISTICAL MODELLING
STATISTICAL PAPERS
STATISTICAL SCIENCE
STATISTICS
STATISTICS and PROBABILITY LETTERS
TEST

Table A.3: Journals in the Applied Statistics category

Applied Statistics
ANNALS OF APPLIED STATISTICS
APPLIED STOCHASTIC MODELS IN BUSINESS AND INDUSTRY
BIOMETRICAL JOURNAL
BIOMETRICS
BIOSTATISTICS
BRITISH JOURNAL OF MATHEMATICAL and STATISTICAL PSYCHOLOGY
CHEMOMETRICS AND INTELLIGENT LABORATORY SYSTEMS
ENVIRONMENTAL AND ECOLOGICAL STATISTICS
ENVIRONMETRICS
IEEE-ACM TRANSACTIONS ON COMPUTATIONAL BIOLOGY AND BIONFORMATICS
INTERNATIONAL JOURNAL OF BIOSTATISTICS
JOURNAL OF AGRICULTURAL BIOLOGICAL AND ENVIRONMENTAL STATISTICS
JOURNAL OF APPLIED STATISTICS
JOURNAL OF BIOPHARMACEUTICAL STATISTICS
JOURNAL OF CHEMOMETRICS
JOURNAL OF COMPUTATIONAL BIOLOGY
JOURNAL OF OFFICIAL STATISTICS
JOURNAL OF QUALITY TECHNOLOGY
JOURNAL OF THE ROYAL STATISTICAL SOCIETY SERIES A STATISTICS IN SOCIETY
JOURNAL OF THE ROYAL STATISTICAL SOCIETY SERIES C APPLIED STATISTICS
MATHEMATICAL POPULATION STUDIES
MULTIVARIATE BEHAVIORAL RESEARCH
OPEN SYSTEMS and INFORMATION DYNAMICS
PHARMACEUTICAL STATISTICS
PROBABILISTIC ENGINEERING MECHANICS
QUALITY ENGINEERING
SORT-STATISTICS AND OPERATIONS RESEARCH TRANSACTIONS
STATISTICAL APPLICATIONS IN GENETICS AND MOLECULAR BIOLOGY
STATISTICAL METHODS IN MEDICAL RESEARCH
STATISTICS IN BIOPHARMACEUTICAL RESEARCH
STATISTICS IN MEDICINE
STOCHASTIC ENVIRONMENTAL RESEARCH AND RISK ASSESSMENT
SURVEY METHODOLOGY
TECHNOMETRICS

Table A.4: Journals in the Computational Statistics category

Computational Statistics
COMMUNICATIONS IN STATISTICS - SIMULATION AND COMPUTATION
COMPUTATIONAL STATISTICS
COMPUTATIONAL STATISTICS and DATA ANALYSIS
JOURNAL OF COMPUTATIONAL AND GRAPHICAL STATISTICS
JOURNAL OF STATISTICAL COMPUTATION AND SIMULATION
JOURNAL OF STATISTICAL SOFTWARE
R JOURNAL
STATA JOURNAL
STATISTICS AND COMPUTING

Table A.5: Journal in the Econometrics and Financial Statistics category

Econometrics and Financial Statistics
ASTIN BULLETIN
ECONOMETRIC REVIEWS
ECONOMETRIC THEORY
ECONOMETRICA
ECONOMETRICS JOURNAL
FINANCE AND STOCHASTICS
INSURANCE MATHEMATICS and ECONOMICS
JOURNAL OF BUSINESS and ECONOMIC STATISTICS
LAW PROBABILITY and RISK
OXFORD BULLETIN OF ECONOMICS AND STATISTICS
QUALITY and QUANTITY
QUALITY TECHNOLOGY AND QUANTITATIVE MANAGEMENT
SCANDINAVIAN ACTUARIAL JOURNAL